

Please type a plus sign (+) inside this box → ☐

PTO/SB/05 (4/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL <small>(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))</small>	Attorney Docket No. 2207/7533
	First Inventor or Application Identifier Gad S. SHEAFFER
	Title BUFFER PRE-LOADING FOR MEMORY SERVICE INTERRUPTIONS
	Express Mail Label No.

APPLICATION ELEMENTS <small>See MPEP chapter 600 concerning utility patent application contents</small>	ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
1. <input checked="" type="checkbox"/> * Fee Transmittal Form (e.g., PTO/SB/17) (Submit an original and a duplicate for fee processing)	5. <input type="checkbox"/> Microfiche Computer Program (Appendix)
2. <input checked="" type="checkbox"/> Specification [Total Pages 20] (preferred arrangement set forth below) <ul style="list-style-type: none">- Descriptive title of the Invention- Cross References to Related Applications- Statement Regarding Fed sponsored R & D- Reference to Microfiche Appendix- Background of the Invention- Brief Summary of the Invention- Brief Description of the Drawings (if filed)- Detailed Description- Claim(s)- Abstract of the Disclosure	6. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) <ul style="list-style-type: none">a. <input type="checkbox"/> Computer Readable Copyb. <input type="checkbox"/> Paper Copy (identical to computer copy)c. <input type="checkbox"/> Statement verifying identity of above copies
3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) [Total Sheets 5]	ACCOMPANYING APPLICATION PARTS
4. Oath or Declaration [Total Pages 3] <ul style="list-style-type: none">a. <input checked="" type="checkbox"/> Newly executed (original or copy)b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) (for continuation/divisional with Box 16 completed)<ul style="list-style-type: none">i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).	7. <input checked="" type="checkbox"/> Assignment Papers (cover sheet & document(s))
* NOTE FOR ITEMS 1 & 13 IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).	
8. <input type="checkbox"/> 37 C.F.R. § 3.73(b) Statement of Power of Attorney (when there is an assignee) <input type="checkbox"/> Attorney	
9. <input type="checkbox"/> English Translation Document (if applicable)	
10. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations	
11. <input type="checkbox"/> Preliminary Amendment	
12. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized)	
13. <input type="checkbox"/> * Small Entity Statement(s) filed in prior application (PTO/SB/09-12) <input type="checkbox"/> Status still proper and desired	
14. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed)	
15. <input type="checkbox"/> Other: _____	

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:

<input type="checkbox"/> Continuation	<input type="checkbox"/> Divisional	<input type="checkbox"/> Continuation-in-part (CIP)	of prior application No. _____
Prior application information		Examiner	Group / Art Unit

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS					
<input type="checkbox"/> Customer Number or Bar Code Label: _____ (Insert Customer No. or Attach bar code label here)			or <input checked="" type="checkbox"/> Correspondence address below		
Name	John C. Altmiller Kenyon & Kenyon				
Address	Suite 700 1500 K Street, NW				
City	Washington	State	DC	Zip Code	20005-1257
Country		Telephone	(202) 220-4200	Fax	(202) 220-4201

Name (Pnnt/Type)	James M. Ross	Registration No. (Attorney/Agent)	42115
Signature	Date		12/20/99

Burden Hour Statement. This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231

BUFFER PRE-LOADING FOR MEMORY SERVICE INTERRUPTIONS

FIELD OF THE INVENTION

Embodiments of the present invention provide a method and apparatus for
5 memory buffer pre-loading. More particularly, embodiments of the present
invention provide methods and apparatus for memory buffer pre-loading for memory
service interruptions.

BACKGROUND OF THE INVENTION

10 In known systems including a memory that handles real time data streams,
memory buffers are implemented in conjunction with the memory to handle the real
time data stream. For example, a real time data stream is handled by a memory of a
computer system when the memory receives and temporarily stores video data. The
video data may be received from graphics processing circuitry (e.g., a graphics
15 coprocessor, a general purpose microprocessor that can perform graphics operations,
etc.), from a video capture device (e.g., a video camera, etc.), from a
communications network (e.g., a cable network, the Internet, etc.), and other known
video data sources. After the video data is temporarily stored by the memory, it may
be read from the memory and sent to a video display, to data storage (e.g., a hard
20 disk, a Digital Versatile Disk (DVD), etc.), over a communications network, etc.

A memory service interruption can result in the memory being unable to
handle real time data streams. Memory buffers may be implemented to handle the
real time data streams during the memory service interruptions. For example, a
memory buffer can be implemented to receive video data. Video data to be written
25 to a memory can be received by a write memory buffer, which temporarily holds

received video data until the video data can be written to the memory. Likewise, video data to be read from a memory may be sent to a read memory buffer, which temporarily holds the read video data until the video data is sent to a video display.

Memory buffers can be sized to handle the memory data demands that can occur during a memory service interruption. For example, a read memory buffer (i.e., read buffer) can have a physical capacity of two kilobytes to hold 256 quadwords (a quadword can correspond to eight bytes). The read buffer also can have a watermark with a set value. The watermark value may equal the number of valid data entries (e.g., quadwords, bytes, etc.) that the buffer should contain prior to a memory service interruption. For example, when 192 data entries are to be read from the memory buffer during the pendency of a worst case memory service interruption, then the watermark of the read buffer can be set at 192 data entries.

FIG. 1 is an illustration of the operation of a known memory buffer. Memory buffer 100 is a read buffer with a physical capacity of holding 144 data entries. Memory buffer 100 is operated with a watermark 101. The watermark 101 has a value corresponding to 100 data entries. Memory buffer operations can be performed to read data entries from the memory into the read buffer such that the read buffer typically is filled to the watermark level (e.g., has at least 100 data entries). When the number of data entries in the read buffer falls below the watermark level (e.g., less than 100 data entries), then high priority, short burst size data requests 102 can be generated to read data entries from the memory into the read buffer. When the number of data entries in the read buffer is equal to or greater than the watermark level (e.g., 100 or more data entries), then low priority, long burst size data requests 103 can be generated to continue to fill the read buffer.

Thus, known memory buffers that handle real time data streams can be filled and allowed to drain to a watermark. The value of the watermark can be set so that the buffer can sustain the worst case latency to memory. Known systems can have both a worst case latency to memory and a regular memory latency. System performance can be disadvantageously affected when memory service interruptions occur, but memory buffer operations are performed to keep the memory buffer at watermark corresponding to the worst case latency to memory. Maintaining a

watermark at a higher level can disadvantageously affect system performance by generating a greater number of high priority, short burst size data requests, which can result in a lower page hit rate to memory.

In view of the foregoing, it can be appreciated that a substantial need exists
5 for a method and apparatus which can more efficiently and effectively manage memory buffer operations.

SUMMARY OF THE INVENTION

Embodiments of the present invention provide methods and apparatus to
10 control the loading of a memory buffer. An advance indication of a memory service interruption can be received by a memory buffer. The memory buffer may have a watermark with a first watermark value. The watermark can be modified to have a second watermark value different from the first watermark value based at least in part on the received advance indication of the memory service interruption.

15

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is an illustration of the operation of a known memory buffer.

FIG. 2 shows an illustration of a system in accordance with an embodiment of the present invention.

20 FIG. 2A shows signals received by an apparatus in accordance with an embodiment of the present invention.

FIG. 3 shows an illustration of the operation of a memory buffer in accordance with an embodiment of the present invention.

FIG. 4 shows a method in accordance with an embodiment of the present
25 invention.

DETAILED DESCRIPTION

Embodiments of methods and systems for memory buffer pre-loading for memory service interruptions are described. In the following description, for
30 purposes of explanation, numerous specific details are set forth to provide a thorough understanding of the present invention. It will be appreciated, however, by

one skilled in the art that the present invention may be practiced without these specific details. In other instances, structures and devices are shown in block diagram form. Furthermore, one skilled in the art can readily appreciate that the specific sequences in which methods are presented and performed are illustrative and it is contemplated that the sequences can be varied and still remain within the spirit and scope of the present invention.

FIG. 2 shows an illustration of a system in accordance with an embodiment of the present invention. Computer 250 may be coupled to a video display 260 via an external computer bus 255. Computer 250 can include a processor 270, such as the Pentium® III processor manufactured by Intel Corp. of Santa Clara, California. In another embodiment, the processor can be an Application Specific Integrated Circuit (ASIC). Computer 250, in one embodiment, includes a memory 220 coupled to the processor 270. The term "coupled" means connected directly or indirectly. The memory 220 is any device adapted to store digital information, such as Dynamic Random Access Memory (DRAM), Rambus® DRAM (RDRAM®), flash memory, a hard disk, an optical digital storage device, or a combination thereof, etc. Rambus® technology is licensed by Rambus Inc. of Mountain View, California, and both Rambus® and RDRAM® are registered trademarks of Rambus Inc. The external computer bus 255 can be an Advanced Graphics Port (AGP) compliant bus.

Processor 270, in one embodiment, includes memory controller 200. In one embodiment, processor 270 is an integrated processor including a graphics module (e.g., a graphics engine) and a memory controller, and is coupled to a Unified Memory Architecture (UMA) memory. In accordance with one embodiment of the present invention, a UMA memory can store a frame buffer (e.g., a section of memory used to store an image to be displayed on a video monitor and/or parts of the image that lie outside the limits of the display, etc.). In one embodiment, the integrated processor can include a processor core, such as a Pentium® processor core. In another embodiment, a processor does not include a memory controller, but is coupled to a memory controller.

Memory controller 200 can be coupled to memory buffer 210 and memory 220. Memory buffer 210 may be coupled to memory 220 to allow data entries to be

transferred between memory buffer 210 and memory 220. When memory buffer 210 is a read buffer, memory controller 200 can generate requests of varied priority (e.g., a high priority request, a low priority request, an intermediate priority request, etc.) to fill the memory buffer 210 with data entries from the memory 220. In one embodiment, the memory buffer 210 handles a real time video data stream that is being read from memory 220 and written to video display 260.

Memory controller 200, in one embodiment, includes a watermark register 201 to store a current value of the watermark of the memory buffer 210. The memory controller 200 can also include an above-watermark burst size register 202 and a below-watermark burst size register 203. When the number of valid data entries in the memory buffer 210 falls below the watermark, a high priority data request to fill the memory buffer 210 can be generated with a below-watermark burst size. The below-watermark burst size can be stored in the below-watermark burst size register 203, which can be a five bit register, for example, that specifies up to 32 different burst sizes. When the number of valid data entries in the memory buffer 210 is above the watermark, a low priority data request to fill the memory buffer 210 can be generated with an above-watermark burst size. The above-watermark burst size can be stored in the above-watermark burst size register 202, which can be a five bit register, for example, that specifies up to 32 different burst sizes.

In one embodiment, the memory 220 includes RDRAM. Operation of RDRAM can include a regular memory latency and a worst case latency to memory. The regular memory latency can be caused by the Rambus memory controller and memory devices. The worst case latency to memory can be caused in part by RDRAM maintenance operations that cause the memory to shut down and not serve memory requests for a period of time. RDRAM maintenance operations can include memory controller temperature calibration, memory device temperature calibration, memory controller current calibration, memory device current calibration, DLL (Delay Loop Line) refresh for device(s) in the Nap mode, memory refresh, etc. In known RDRAM implementations, RDRAM memory maintenance operations may occur every 100 milliseconds and have durations of up to 2.5 microseconds.

5 The memory buffer 210 can typically be operated based on the regular memory latency. During such operation, the watermark of memory buffer 210 can have a regular watermark value, the above-watermark burst size can have a regular above-watermark burst size value, and the below-watermark burst size can have a regular below-watermark burst size value. In one embodiment, the regular watermark value is stored in the regular watermark value register 205, the regular above-watermark burst size value is stored in the regular, above-watermark burst size value register 207a, and regular below-watermark burst size value is stored in the regular, below-watermark burst size value register 207b. When the memory
10 buffer 210 is operating based on the regular memory latency, the watermark register 201 can contain the regular watermark value, the above-watermark burst size register 202 can contain the regular above-watermark burst size value, and the below-watermark burst size register 203 can contain the regular below-watermark burst size value.

15 The memory controller 200, in one embodiment, receives a memory service interruption advance indication signal that can provide advance warning that a memory service interruption having a worst-case latency to memory will occur. The amount of time the advance warning precedes the memory service interruption having a worst-case latency to memory can be varied based on system requirements, performance parameters, etc. For example, the memory service interruption
20 indication signal can provide an advance indication that an RDRAM maintenance operation will be performed, and the timing of the advance indication can be based on a pre-maintenance period length value stored in pre-maintenance period length register 204. Subsequent to receiving the indication that an RDRAM maintenance operation will be performed, the memory controller 200 can modify the watermark value, the above-watermark burst size value, and the below-watermark burst size value of memory buffer 210. The pre-maintenance watermark value stored in the pre-maintenance watermark value register 206 can be written to the watermark
25 register 201, and the pre-maintenance, above-watermark burst size value stored in the pre-maintenance, above-watermark burst size value register 208a can be written
30 to the above-watermark burst size register 202, and the pre-maintenance, below-

watermark burst size value stored in the pre-maintenance, below-watermark burst size value register 208b can be written to the below-watermark burst size register 203. In another embodiment, the above-watermark burst size and/or the below-watermark burst-size can be constant and hardwired, rather than variable and stored in a register.

Subsequent to modifying the watermark from the regular value to the pre-maintenance value, the memory controller 200 can generate high priority, below-watermark burst size requests to fill the memory buffer 210 to the pre-maintenance watermark level when the level of memory buffer 210 is below the pre-maintenance watermark level. When the memory buffer 210 is filled above the pre-maintenance watermark level, the memory controller 200 can generate low priority, pre-maintenance above-watermark burst size requests to fill the memory buffer 210.

The memory controller 200, in one embodiment, receives a memory service interruption status signal that can provide status information about a memory service interruption having a worst-case latency to memory. For example, the memory service interruption status signal can indicate when an RDRAM maintenance operation is being performed. In one embodiment, the memory service interruption status signal goes high at the onset of an RDRAM maintenance operation and goes low at the termination of an RDRAM maintenance operation. Subsequent to receiving an indication that the RDRAM maintenance operation has terminated, the memory controller 200 can modify the watermark value, the above-watermark burst size value, and the below-watermark burst size value to their respective regular values. The regular watermark value stored in the regular watermark value register 205 can be written to the watermark register 201, the regular, above-watermark burst size value stored in the regular, above-watermark burst size value register 207a can be written to the above-watermark burst size register 202, and the regular, below-watermark burst size value stored in the regular, below-watermark burst size value register 207b can be written to the below-watermark burst size register 203.

FIG. 2A shows signals received by an apparatus in accordance with an embodiment of the present invention. Memory service interruption status signal 280 can go high during an RDRAM maintenance operation. In other embodiments of

the present invention, the memory service interruption status signal can indicate the status of another memory maintenance operation, a Rambus dynamic random access memory maintenance operation, a memory device temperature calibration, a memory device current calibration, a memory refresh, etc. In one embodiment, the

5 memory service interruption status signal 280 is generated by a timer and can cause initiation of the RDRAM maintenance operation. The transition of the memory service interruption status signal 280 from low to high can indicate the onset of an RDRAM maintenance operation. The transition of the memory service interruption status signal 280 from high to low can indicate the termination of an RDRAM

10 maintenance operation. Status indications 281 and 282 each correspond to the pendency of an RDRAM maintenance operation. In one embodiment, status indications 281 and 282 are 100 milliseconds apart and each have a duration of 2.5 milliseconds. In other embodiments, the timing and duration of status indications can vary based on system parameters, performance requirements, etc.

15 Memory service interruption advance indication signal 285 can go high prior to an RDRAM maintenance operation to provide an advance indication of an upcoming RDRAM maintenance operation. In other embodiments of the present invention, a memory service interruption indication signal can include an indication of a memory maintenance operation, an indication of a Rambus dynamic random

20 access memory maintenance operation, an indication of a memory device temperature calibration, an indication of a memory device current calibration, an indication of a memory refresh, etc. In one embodiment, the memory service interruption advance indication signal 285 is generated by a timer. Advance indication 286 provides an advance warning of the RDRAM maintenance operation corresponding to status indication 281. Likewise, advance indication 287 provides

25 an advance warning of the RDRAM maintenance operation period corresponding to status indication 282. In one embodiment, advance indications 286 and 287 are 100 milliseconds apart and can have a duration of 5 milliseconds. In another embodiment, advance indications 286 and 287 stay high until the end of the memory

30 service interruption. Advance indications 286 and 287 respectively can precede status indications 281 and 282 by approximately 500 nanoseconds. In other

embodiments, the timing and duration of advance indications can vary based on system parameters, performance requirements, etc.

Memory buffer status timeline 290 includes timeline segments 291, 292, 293, 294, and 295. During timeline segment 291, the memory buffer 210 is operated with the regular watermark and the regular burst sizes. Timeline segment 291 ends, and timeline segment 292 begins, with the onset of advance indication 286. In response to advance indication 286, the memory controller 200 modifies the watermark value and the burst size values and operates the memory buffer 210 during timeline segment 292 with the pre-maintenance watermark and the pre-maintenance burst sizes. Timeline segment 292 ends, and timeline segment 293 begins, when status indication 281 terminates. In response to the termination of status indication 281, the memory controller 200 modifies the watermark value and the burst size values and operates the memory buffer 210 during timeline segment 293 with the regular watermark and the regular burst sizes. In another embodiment where the advance indications 286, 287 of memory service interruption advance indication signal 285 stay high until the end of the memory service interruption, when the advance indications 286, 287 end (e.g., the memory service interruption advance indication signal 285 goes low, etc.), the memory controller 200 can modify the watermark value and the burst size values and operates the memory buffer 210 during timeline segment 293 with the regular watermark and the regular burst sizes. The operation of the memory buffer 210 during timeline period 294 corresponds to the operation of the memory buffer during timeline period 292. Likewise, the operation of the memory buffer 210 during timeline period 295 corresponds to the operation of the memory buffer during timeline period 293.

FIG. 3 shows an illustration of the operation of a memory buffer in accordance with an embodiment of the present invention. Memory buffer 300 can be a read buffer that buffers a real time video stream that is stored in a memory and being sent to a video display (e.g., video display 260 of FIG. 2). Memory buffer 300 can be filled with 96 data entries, each data entry being a quadword of eight bytes. The memory buffer can typically be operated to handle a memory service interruption having a regular memory latency. For example, in one embodiment, the

memory buffer 300 can be filled to the regular watermark 310 level (i.e., be filled with 30 data entries) to handle a memory service interruption having a regular memory latency. In such a regular mode of operation, when the level of the memory buffer falls below the regular watermark 310, then high priority, regular below-watermark, burst size data requests (illustrated by label 311) are generated by a memory controller (e.g., memory controller 200 of FIG. 2) to fill the memory buffer. In one embodiment, the regular below-watermark burst size is 8 data entries. In other embodiments, the regular below-watermark burst size can include a different number of data entries based on system parameters, performance requirements, etc.

When the level of the memory buffer 300 is above the regular watermark 310 during the regular mode of operation, then low priority, regular above-watermark burst size data requests (illustrated by label 312) are generated by a memory controller to fill the memory buffer. The size of the buffer can be computed so that there is enough space between the top of the buffer and the regular watermark to work most of the time using low priority, regular above-watermark burst size data requests. In one embodiment, when the regular memory latency is 500 nanoseconds and the regular watermark level is 30 data entries, then the memory buffer size is 96 data entries and the regular above-watermark burst size is 24 data entries or the number of empty buffer entries. In other embodiments, the buffer size and burst sizes can vary based on system parameters, performance requirements, etc.

To handle the longer memory latency associated with memory service interruption having a worst case latency to memory (e.g., an RDRAM maintenance operation), the value of the watermark and the above-watermark and below-watermark burst sizes can be changed. In one embodiment, 0.5 milliseconds (or another value stored in pre-maintenance period length register 204 by a driver) before a worst case latency to memory service interruption, operation of the memory buffer can be modified to handle the worst case latency to memory. For example, in one embodiment, the memory buffer 300 should be filled to the worst case latency watermark 320 level (e.g., be filled with 78 data entries) to handle a memory service interruption having a worst case latency to memory of 2 milliseconds. Thus, the value of the memory buffer watermark level is changed from the regular watermark

level (e.g., 30 data entries) to the worst case latency watermark level (e.g., 78 data entries). The increased level of the watermark (e.g., an additional 48 data entries) can cover the additional latency caused by a worst case latency to memory service interruption.

5 In such a worst case latency mode of operation, when the level of the memory buffer falls below the worst case latency watermark 320, then high priority, worst case latency below-watermark burst size data requests (illustrated by label 321) are generated by a memory controller to fill the memory buffer. When the level of the memory buffer 300 is above the worst case latency watermark 320
10 during the worst case latency mode of operation, then low priority, worst case latency above-watermark burst size data requests (illustrated by label 322) are generated by a memory controller to fill the memory buffer. In one embodiment, the worst case latency, above-watermark burst size equals the difference between the size of the memory buffer and the worst case latency watermark level. For example,
15 in the embodiment illustrated in FIG. 3, the worst case latency, above-watermark burst size can be 18 data entries (i.e., 96 data entries - 78 data entries = 18 data entries). Such a burst size can have the effect of keeping the buffer as full as possible.

 The embodiment of the present invention illustrated in FIG. 3 is a read
20 memory buffer that is operated to be filled with a certain number of data entries to handle memory service operations (e.g., a memory service interruption). Other embodiments of the present invention can encompass control of other memory buffers, such as write memory buffers. In a write memory buffer, a watermark can be set to operate the write memory buffer such that a certain number of memory
25 locations are available to receive data entries to be written to the memory. For example, a write memory buffer may receive a real time video stream from a video capture device (e.g., a video camera), a video communications device (e.g., a cable television network connection), or another known source of video data. During a memory service interruption, the write memory buffer should have adequate
30 capacity to handle the real time data stream that is to be written to the memory. Thus, in one embodiment of the present invention, a write memory buffer is

operated to handle a regular memory latency. An indication of a memory service interruption having a different latency to memory is received, and the operation of the write memory buffer is modified (e.g., watermark value modified, burst size values modified, etc.) to handle the different latency to memory.

5 FIG. 4 shows a method in accordance with an embodiment of the present invention. In one embodiment of the present invention, a method can be implemented by instructions to be executed by a processor. At block 400, a memory buffer (e.g., a memory buffer 210 of FIG. 2) is operating in a regular memory latency mode. The memory buffer's watermark has a regular watermark value and
10 the below-watermark burst size has a regular below-watermark burst size value. At block 401, a memory controller (e.g., a memory controller 200 of FIG. 2) determines whether an indication of a memory service interruption having a worst case latency to memory has been received. If no such indication has been received, then there is no change to the memory buffer control parameters (e.g., the watermark, the below-
15 watermark burst size, etc.), (block 402), and the buffer continues to operate in the regular memory latency mode until such an indication has been received. If an indication of a memory service interruption having a worst case latency to memory has been received, then the watermark value is changed to the worst case latency watermark value and the below-watermark burst size is changed to the worst case
20 latency burst size value (block 403). Subsequent to execution of block 403, the memory buffer is operating in worst case latency mode.

At block 404, the memory controller determines whether an indication of the termination of the memory service interruption having a worst case latency to memory has been received. If no such indication has been received, then there is no
25 change to the memory buffer control parameters (e.g., the watermark, the below-watermark burst size, etc.), (block 405), and the buffer continues to operate in the worst case latency mode until such an indication has been received. If an indication of the termination of a memory service interruption having a worst case latency to memory has been received, then the watermark value is restored to the regular
30 watermark value and the below-watermark burst size restored to the regular below-watermark burst size value (block 406). Subsequent to execution of block 406, the

memory buffer is again operating in regular memory latency mode, and at block 407 execution loops back to the beginning.

In accordance with one embodiment of the present invention, instructions adapted to be executed to perform a method in accordance with an embodiment of the present invention are stored on a computer-readable medium. The computer-readable medium can be any device adapted to store digital information. The computer-readable medium can be distributed to a user that has a processor suitable for executing instructions adapted to be executed. The term "adapted to be executed" is meant to encompass any instructions that are ready to be executed in their present form (e.g., machine code) by a processor, or require further manipulation (e.g., compilation, decryption, or provided with an access code, etc.) to be ready to be executed by a processor.

Embodiments of the present invention advantageously allow memory buffer pre-loading for memory service interruptions. In an embodiment having regular memory latency service interruptions and worst case latency to memory service interruptions, memory buffer watermark and/or burst size values can be modified to provide efficient system performance. Embodiments of the present invention may provide for an increase in the proportion of low priority, longer burst size data requests, which can result in greater page hit rates to memory.

In the foregoing detailed description, apparatus and methods in accordance with embodiments of the present invention have been described with reference to specific exemplary embodiments. Accordingly, the present specification and figures are to be regarded as illustrative rather than restrictive.

What is claimed is:

- 1 1. A method to control the loading of a memory buffer, the memory
2 buffer having a watermark with a first watermark value, the method comprising:
3 receiving an advance indication of a memory service interruption;
4 and
5 based at least in part on the received advance indication of the
6 memory service interruption, modifying the watermark to have a second
7 watermark value different from the first watermark value.
- 1 2. The method of claim 1, wherein the memory buffer having a
2 watermark has a below-watermark burst size with a first burst size value, the method
3 further comprising:
4 based at least in part on the received advance indication of the
5 memory service interruption, modifying the below-watermark burst size to
6 have a second burst size value different from the first burst size value.
- 1 3. The method of claim 1, wherein the indication of a memory service
2 interruption includes an advance indication of a memory service interruption having
3 a worst case latency to memory.
- 1 4. The method of claim 1, wherein the second watermark value is
2 greater than the first watermark value.
- 1 5. The method of claim 2 wherein the second burst size value is less
2 than the first burst size value.
- 1 6. The method of claim 2 wherein the second burst size value
2 corresponds to difference between the number of data entries in the memory buffer
3 and the second watermark value.

1 7. The method of claim 1, the method further comprising
2 receiving an indication of the termination of the memory service
3 interruption; and
4 based at least in part on the received indication of the termination of
5 the memory service interruption, modifying the watermark to have a third
6 watermark value different from the second watermark value.

1 8. The method of claim 7, wherein the third watermark value equals the
2 first watermark value.

1 9. The method of claim 2, the method further comprising
2 receiving an indication of the termination of the memory service
3 interruption;
4 based at least in part on the received indication of the termination of
5 the memory service interruption, modifying the watermark to have a third
6 value different from the second watermark value; and
7 based at least in part on the received indication of the termination of
8 the memory service interruption, modifying the below-watermark burst size
9 to have a third burst size value different from the second burst size value.

1 10. The method of claim 9, wherein the third watermark value equals the
2 first watermark value, and wherein the third burst size value equals the first burst
3 size value.

1 11. A method to control the loading of a memory buffer, the memory
2 buffer having a watermark with a first watermark value, the method comprising:
3 modifying the watermark to have a second watermark value prior to
4 the occurrence of a memory service interruption, the second watermark value
5 being different than the first watermark value; and
6 modifying the watermark to have a third watermark value subsequent
7 to the occurrence of the memory service interruption, the third watermark
8 value being different than the second watermark value.

1 12. The method of claim 11, wherein the third watermark value equals
2 the first watermark value.

1 13. The method of claim 11, wherein the memory buffer having a
2 watermark with a first watermark value has an below-watermark burst size with a
3 first burst size value, the method further comprising:
4 modifying the below-watermark burst-size to have a second burst size
5 value prior to the occurrence of a memory service interruption, the second
6 burst size value being different than the first burst size value; and
7 modifying the below-watermark burst size to have a third burst size
8 value subsequent to the occurrence of the memory service interruption, the
9 third burst size value being different than the second burst size value.

1 14. The method of claim 13, wherein the first burst size value is equal to
2 the third burst size value.

1 15. An apparatus to control the loading of a memory buffer, comprising:
2 a memory buffer; and
3 a memory controller, coupled to said memory buffer, including
4 a watermark register;
5 a first register, coupled to said watermark register, to store a
6 first watermark value; and
7 a second register, coupled to said watermark register, to store
8 a second watermark value.

1 16. The apparatus of claim 15, wherein the memory controller includes:
2 a below-watermark burst size register;
3 a third register, coupled to said below-watermark burst size
4 register, to store a first below-watermark burst size value; and
5 a fourth register, coupled to said below-watermark burst size
6 register, to store a second below-watermark burst size value.

1 17. The apparatus of claim 15, wherein the memory controller is to:
2 receive an advance indication of a memory service
3 interruption;
4 read the second watermark value from said second register
5 based at least in part on the received advance indication of a memory
6 service interruption; and
7 store the second watermark value in said watermark register.

1 18. The apparatus of claim 16, wherein the memory controller is to
2 receive an advance indication of a memory service
3 interruption;
4 read the second watermark value from said second register
5 based at least in part on the received advance indication of a memory
6 service interruption; and
7 store the second watermark value in said watermark register;
8 read the second below-watermark burst size value from said
9 fourth register based at least in part on the received advanced
10 indication of a memory service interruption; and
11 store the second below-watermark burst size value in said
12 below-watermark burst size register.

1 19. A system to process video signals, the system comprising:
2 a processor;
3 a memory, coupled to said processor;
4 a memory buffer, coupled to said memory; and
5 a memory controller, coupled to said memory buffer, including
6 a watermark register;
7 a first register, coupled to said watermark register, to store a
8 first watermark value; and
9 a second register, coupled to said watermark register, to store
10 a second watermark value.

1 20. The system of claim 19, wherein the memory controller is to:
2 receive an advance indication of a memory service
3 interruption;
4 read the second watermark value from said second register
5 based at least in part on the received advance indication of a memory
6 service interruption; and
7 store the second watermark value in said watermark register.

1 21. A computer-readable medium storing a plurality of instructions to be
2 executed by a processor to control a memory buffer having a watermark with a first
3 watermark value and a below-watermark burst size with a first burst size value, said
4 plurality of instructions comprising instructions to:
5 receive an advance indication of a memory service interruption; and
6 based at least in part on the received advance indication of the
7 memory service interruption, modify the watermark to have a second
8 watermark value different from the first watermark value.

1 22. The computer-readable medium of claim 21, further comprising
2 instructions to:
3 based at least in part on the received advance indication of the
4 memory service interruption, modify the below-watermark burst size to have
5 a second burst size value different from the first burst size value.

ABSTRACT

Methods and apparatus control the loading of a memory buffer. The memory buffer may have a watermark with a first watermark value and can receive an advance indication of a memory service interruption. Based at least in part on the received advance indication of the memory service interruption, the watermark can be modified to have a second watermark value different from the first watermark value.

PRIOR ART

FIG. 1

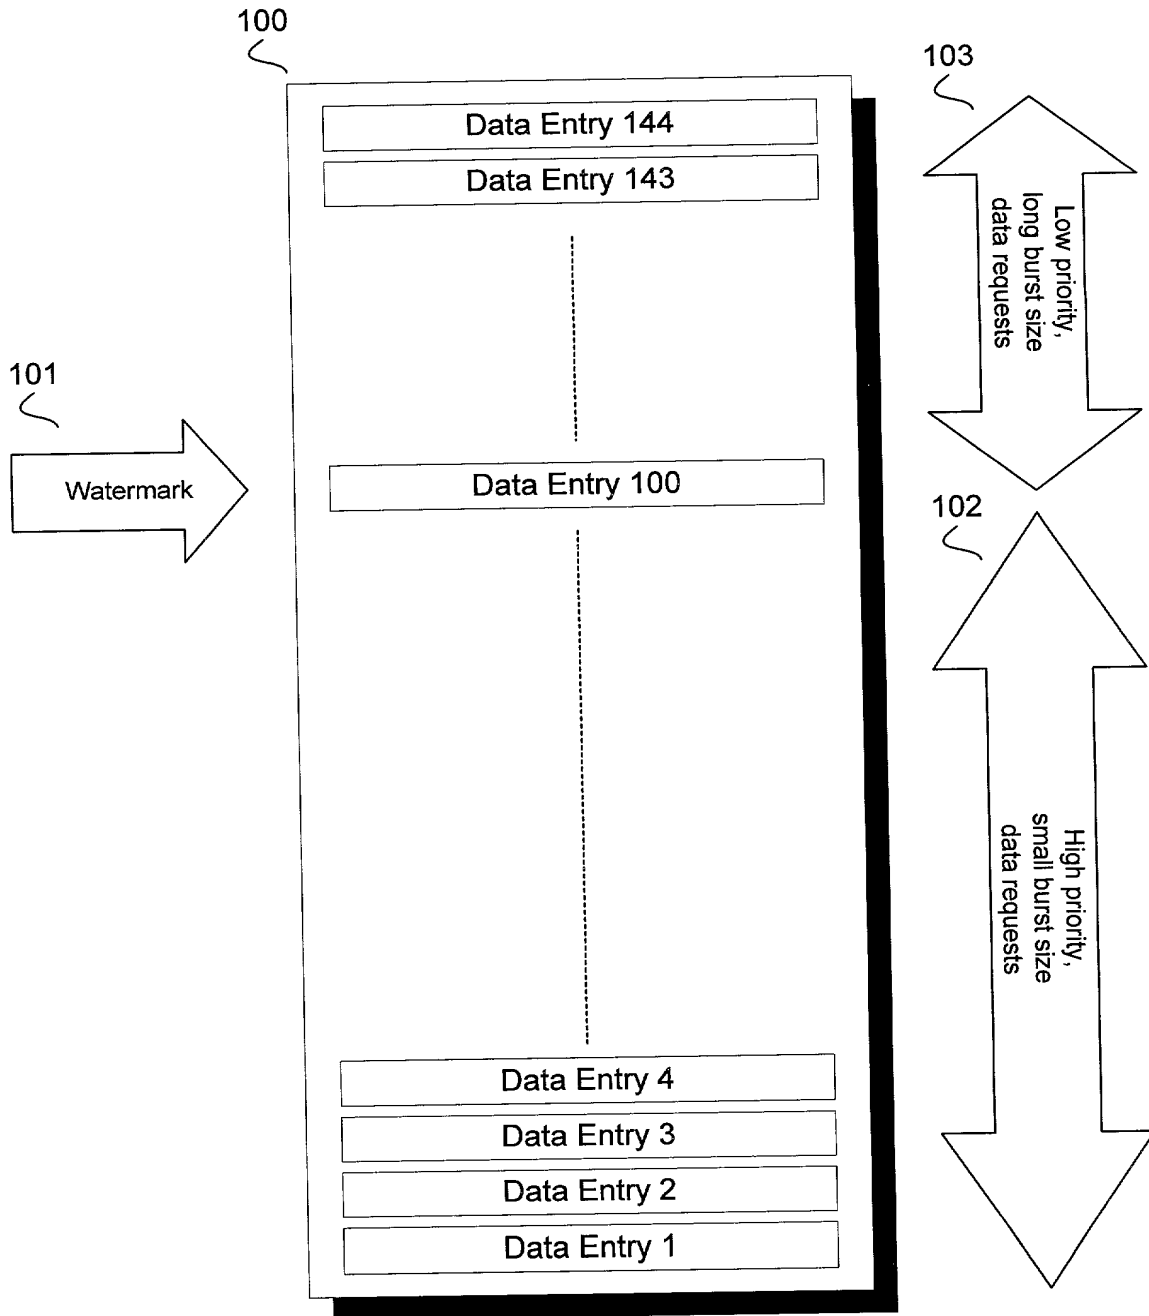


FIG. 2




FIG. 2A




FIG. 3

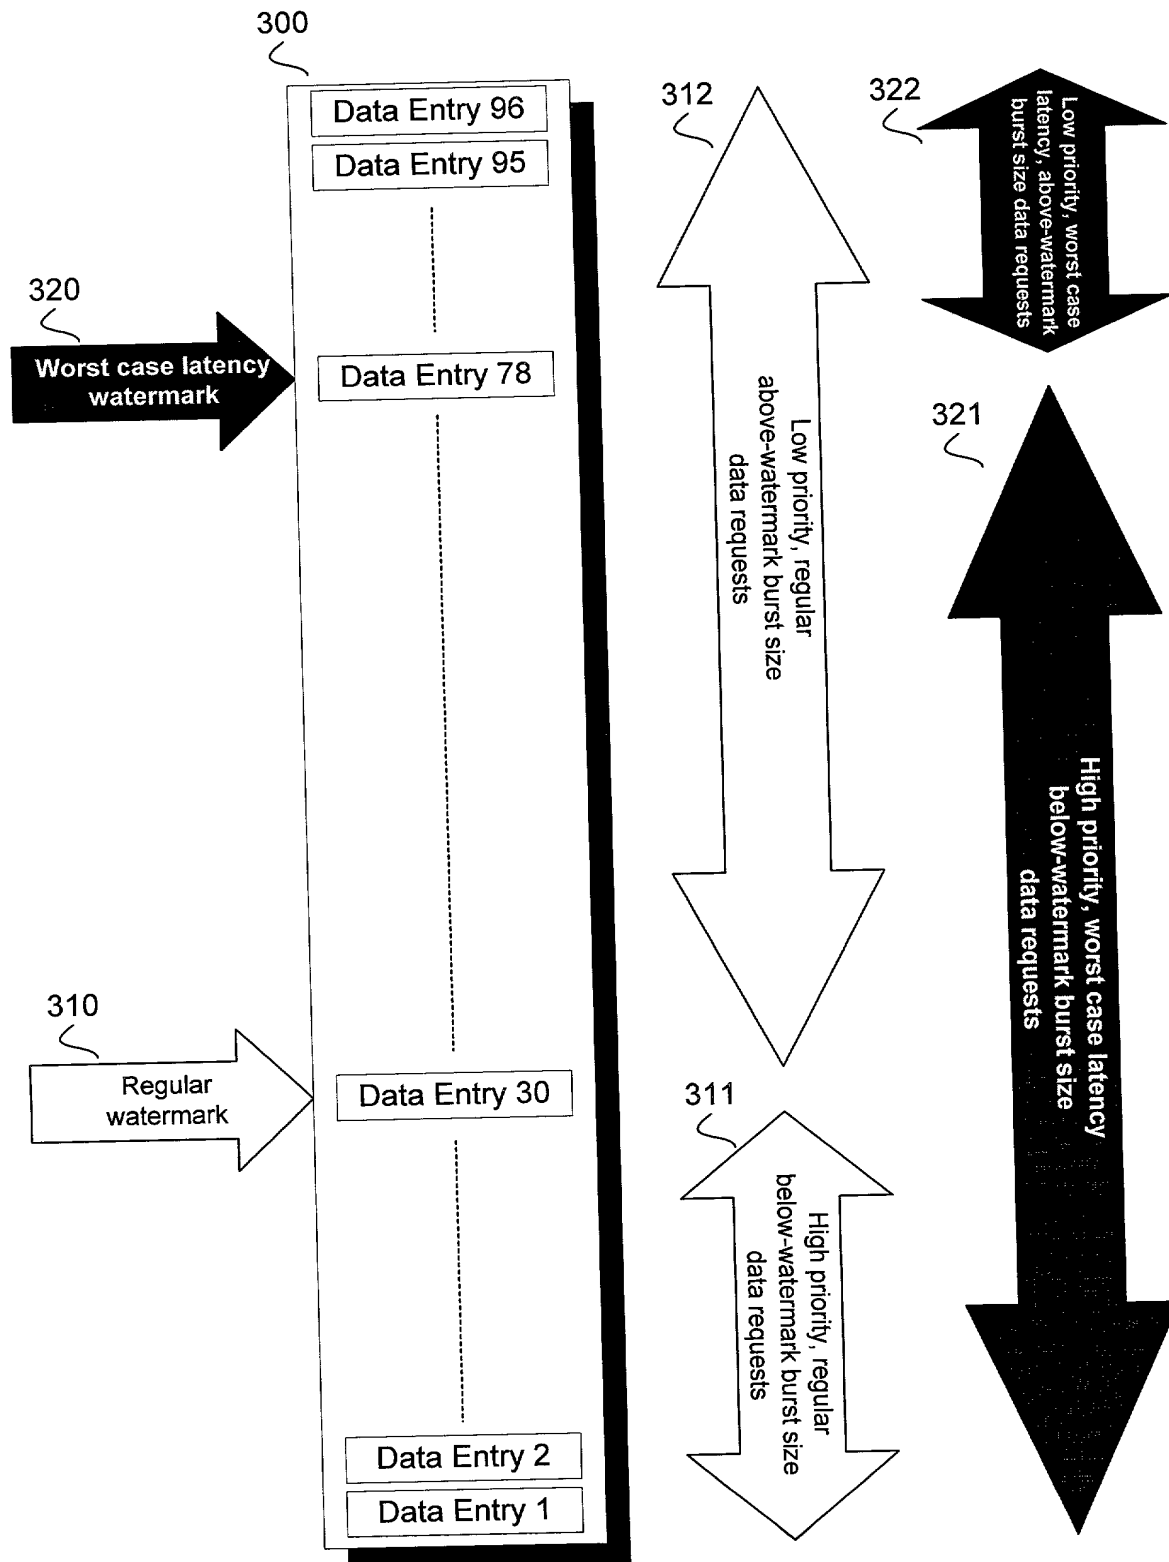
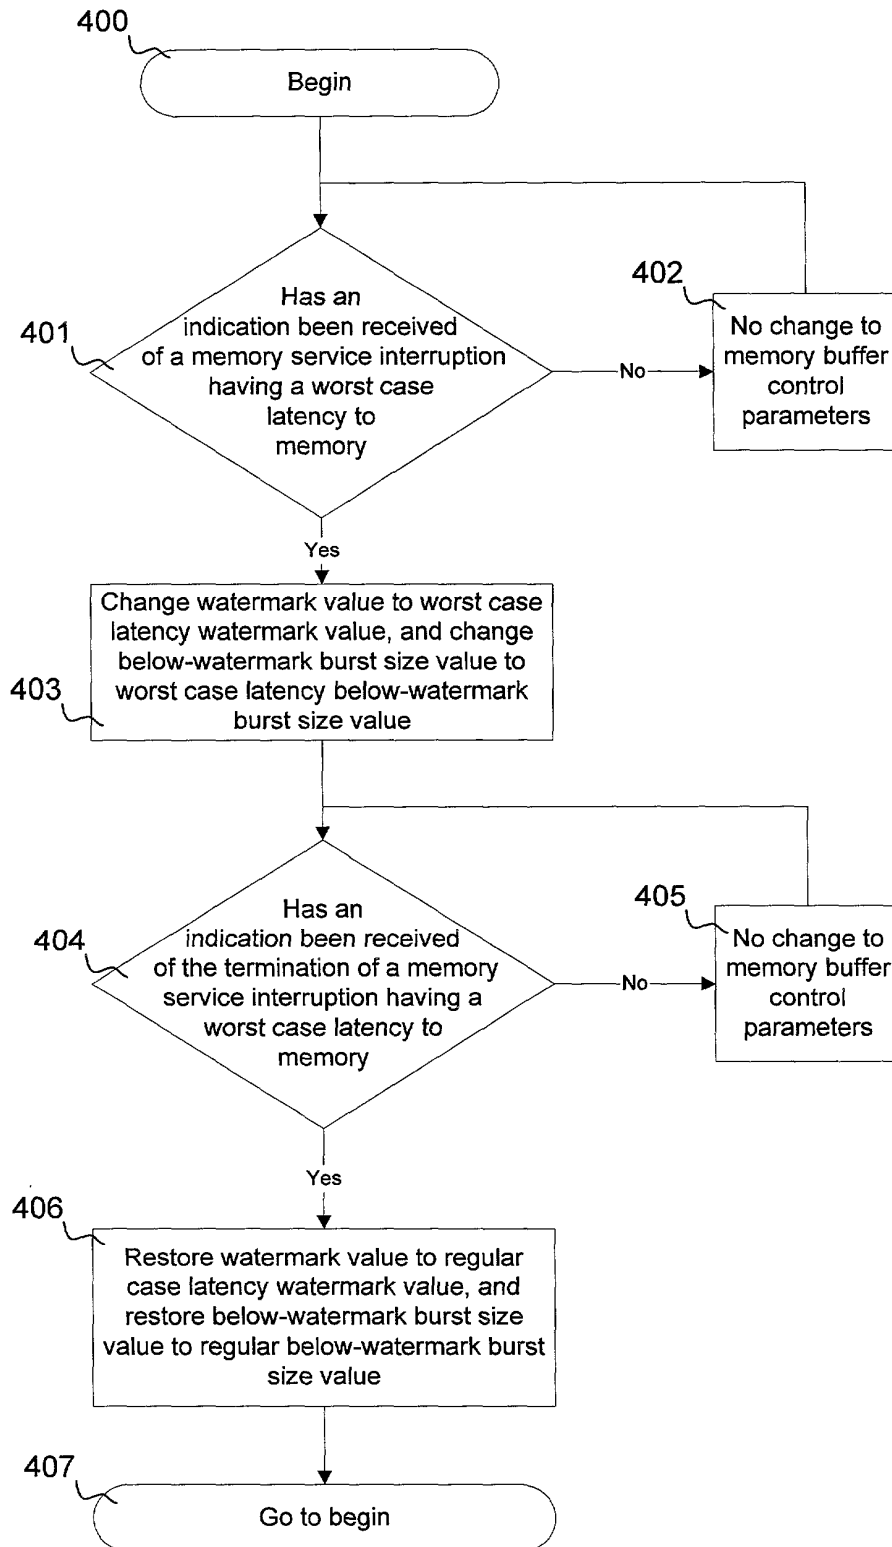


FIG. 4



PATENT
Docket No. 22077633

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address, and citizenship are as stated below next to my name.

I believe I am the original, first, and sole inventor of the subject matter that is claimed and for which a patent is sought on the invention entitled

BUFFER PRE-LOADING FOR MEMORY SERVICE INTERRUPTIONS

the specification of which

___ is attached hereto.

☒ filed herewith.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claim(s). I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, § 1.56(a), a copy of which is attached.

PRIOR FOREIGN APPLICATION(S)

I hereby claim foreign priority benefits under Title 35, United States Code, § 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed.

APPLICATION NUMBER	COUNTRY	FILING DATE (day, month, year)	PRIORITY CLAIMED	
None			Yes	No

PRIOR UNITED STATES APPLICATION(S)

I hereby claim the benefit under Title 35, United States Code, §§ 119-120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56(a) which occurred between the filing date of the prior application and the national or PCT International filing date of this application.

APPLICATION NUMBER	FILING DATE (day, month, year)	STATUS (i.e. Patented, Pending, Abandoned)
None		

POWER OF ATTORNEY: I hereby appoint:


Paul H. Haller (Reg. No. 21,074); John C. Altmiller (Reg. No. 25,951); Felix L. D'Arenzo, Jr. (Reg. No. 27,831); Shawn W. O'Dowd (Reg. No. 34,667); James M. Ross (Reg. No. 42,115) of KENYON & KENYON with offices located at 1500 K Street, N.W., Washington, D.C. 20005, telephone (202) 220-4200, and James E. Jacobson, Jr. (Reg. No. 31,626); Thomas O. Reynolds (Reg. No. 32,480); Raymond J. Werner (Reg. No. 34,752); Richard C. Calderwood (Reg. No. 35,168); Joseph R. Bond (Reg. No. 30,458); Naomi Obinola (Reg. No. 39,320) of INTFL CORPORATION my attorneys with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

03466647 221299

SEND CORRESPONDENCE, AND DIRECT TELEPHONE CALLS TO:

John C. Altmillor
KENYON & KENYON
Suite 700
1500 K Street, N.W.
Washington, D.C. 20005
(202) 220-4200 (phone)
(202) 220-4201 (facsimile)

I hereby declare that all statements made herein of my own knowledge are true and all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under § 1001 of Title 18 of the United States Code and that such willful statements may jeopardize the validity of the application or any patent issuing thereon.

FULL NAME OF INVENTOR	FAMILY NAME SHEAFFER	FIRST GIVEN NAME Gad	SECOND GIVEN NAME S.
RESIDENCE & CITIZENSHIP	CITY Haifa	STATE OR FOREIGN COUNTRY ISRAEL	COUNTRY OF CITIZENSHIP ISRAEL
POST OFFICE ADDRESS	POST OFFICE ADDRESS 33a Einstein Street	CITY Haifa 34604	STATE & ZIP CODE/COUNTRY ISRAEL
Signature 	Date 19 DEC 1999		

SECRET 4369450

Title 37, Code of Federal Regulations. Section 1.56
Duty to Disclose Information Material to Patentability

- (a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of an evaluation of the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability if it is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.87(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:
- (1) Prior art cited in search reports of a foreign patent office in a counterpart application; and
 - (2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein disclosed to the Office.
- (b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and
- (1) It establishes, by itself or in combination with other information, a *prima facie* case of unpatentability of a claim; or
 - (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.
- A *prima facie* case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.
- (c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:
- (1) Each inventor named in the application;
 - (2) Each attorney or agent who prepares or prosecutes the application; and
 - (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.
- (d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.